

IN THE CLAIMS:

Please cancel Claims 17-38, 42-44 without prejudice

- 1 1. (Original) In a file server having a storage operating system, a method for managing
2 storage of data in a plurality of storage devices, each storage device comprising a plural-
3 ity of blocks for storing data, the method comprising the steps of:
- 4 generating block layout information in a file system layer of the storage operating
5 system by determining which blocks within the plurality of blocks are allocated for stor-
6 ing data and which are unallocated;
- 7 transferring the block layout information from the file system layer to a RAID
8 layer of the storage operating system; and
- 9 responsive to the block layout information, the RAID layer controlling the execu-
10 tion of I/O operations by identifying blocks within the plurality of blocks for use by the
11 I/O operations so as to substantially maximize chain lengths of reads for calculation of
12 parity and of writes for data write operations thereof while substantially minimizing cost
13 of calculation of parity,
- 14 wherein the block identification comprises the steps of
- 15 examining all the I/O operations,
- 16 selecting a method for parity calculation which substantially minimizes cost of the
17 parity calculation for the I/O operations, and
- 18 responsive to the block layout information and the parity calculation method se-
19 lection, identifying the blocks within the plurality of blocks for use by the I/O operations.

- 1 2. (Original) A method for managing storage of data in a plurality of storage devices,
2 each comprising a plurality of storage blocks, the method comprising the steps of:

3 generating block layout information; and
4 in response to the block layout information, controlling the execution of an I/O op-
5 eration by identifying storage blocks for use by the I/O operation so as to substantially
6 minimize cost of calculation of error correction parameters across a stripe.

1 3. (Original) The method of claim 2 wherein the calculation of error correction parame-
2 ters comprises the calculation of parity.

1 4. (Original) The method of claim 3 wherein the calculation of parity comprises selecting
2 a parity calculation operation from a group consisting of a subtraction method and a par-
3 ity re-calculation method.

1 5. (Original) The method of claim 2 wherein the identification of storage blocks for use
2 in the I/O operation substantially maximizes the chain length by substantially maximizing
3 the number of blocks having a contiguous physical layout on a storage device.

1 6. (Original) The method of claim 2, further comprising the step of identifying storage
2 blocks for use in the I/O operation so as to substantially maximize the chain length by
3 substantially maximizing the number of blocks having sequential VBN's associated with
4 the storage blocks.

1 7. (Original) The method of claim 2, further comprising the step of identifying storage
2 blocks for use in the I/O operation so as to substantially maximize the chain length by
3 substantially maximizing the locality of the blocks of a storage device.

1 8. (Currently Amended) The method of claim 2 wherein the execution controlling step
2 comprises the steps of:

3 examining blocks to which data is to be written prior to write operations;
4 selecting one of a plurality of parity calculation methodologies including,
5 a first methodology comprising minimizing the number of blocks read, and
6 a second methodology comprising maximizing chain lengths of blocks read for
7 the parity calculation.

1 9. (Original) The method of claim 8, wherein the execution controlling step further com-
2 prises the steps of:

3 implementing the selection responsive to the block layout information; and
4 wherein, if the selection constitutes substantially minimizing the number of read
5 blocks,
6 determining on a stripe-by-stripe basis whether to calculate parity based on a sub-
7 traction method or a recalculation method,
8 performing any appropriate read operations to support the method selected, and
9 calculating parity responsive to the read blocks and the data to be written; and
10 wherein, if the selection constitutes substantially maximizing chain lengths of blocks
11 read,
12 deciding which storage blocks to read to substantially maximize chain length while
13 substantially minimizing the number of storage blocks read to support either a subtraction
14 method or a recalculation method,
15 performing read operations on the blocks to be read, and
16 calculating parity responsive to the read blocks and the data to be written.

1 10. (Original) The method of claim 2, wherein the identification of storage blocks is
2 based at least in part on an available resource.

1 11. (Original) The method of claim 2 further comprising transmitting the block layout
2 information from a file system layer to a RAID layer.

1 12. (Original) The method of claim 2 wherein the generating step further comprises:
2 making a first determination as to whether a storage block is unallocated;
3 making a second determination as to a current implementation of the plurality of
4 storage devices; and
5 generating the block layout information based at least in part on the first and the sec-
6 ond determinations.

1 13. (Original) The method of claim 2, wherein the I/O operation is one of a plurality of
2 I/O operations and one of the plurality of I/O operations is a read operation.

1
2 14. (Original) The method of claim 2, wherein the chain length is a chain length of a read
3 operation for calculation of parity.

1 15. (Original) The method of claim 2, wherein the chain length is a chain length for a
2 write operation for the data.

1 16. (Original) A method for managing storage of data in a storage system comprising a
2 plurality of storage devices each comprising a plurality of storage blocks, the method
3 comprising writing data to predetermined storage blocks across a plurality of stripes and
4 to predetermined storage blocks within each storage device so as to substantially maxi-
5 mize chain length of storage blocks within each storage device while substantially mini-

6 mizing cost of calculation of error correction parameters across each stripe of the plural-
7 ity of stripes.

17. - 38. (Cancelled)

1 39. (Original) A storage system comprising:
2 a plurality of storage devices each comprising:
3 a plurality of storage blocks; and
4 a storage manager in communication with the plurality of storage devices, the stor-
5 age manager writing data to predetermined storage blocks across a plurality of stripes and
6 to predetermined storage blocks within each storage device so as to substantially maxi-
7 mize chain length of storage blocks within a storage device while substantially minimiz-
8 ing the calculation of error correction parameters across each stripe of the plurality of
9 stripes.

1 40. (Original) A system for managing the storage of data, the system comprising:
2 a plurality of storage devices each having a plurality of storage blocks;
3 a storage device manager in communication with the plurality of storage blocks;
4 a block layout information generator in communication with the storage device man-
5 ager and the plurality of storage blocks; and
6 an error correction parameter calculator in communication with the plurality of stor-
7 age blocks and the storage device manager,
8 wherein the storage device manager, in response to the block layout information
9 from the block layout information generator, controls the execution of an I/O operation
10 by identifying storage blocks for use by the I/O operation so as to substantially maximize
11 chain length within the storage device while substantially minimizing the calculation by
12 the error correction parameter calculator of error correction parameters across a stripe.

1 41. (Original) A method for managing storage of data in storage blocks, the method
2 comprising the steps of:
3 generating block layout information;
4 dynamically determining a first number of error correction calculations;
5 dynamically determining a second number corresponding to a chain length; and
6 in response to the block layout information, controlling the execution of an I/O op-
7 eration by identifying storage blocks for use by the I/O operation so as to have a chain
8 length of the second number within a storage device while performing the first number of
9 calculations of error correction parameters across a stripe.

42. - 44. (Cancelled)